Buyers and Sellers in Affine Logic

Jason Reed

January 23, 2016

1 Introduction

Everyone loves that Peyton-Jones/Eber/Seward paper [PJES00], and indeed to a PL person it's lovely: just think real hard, write down the right combinators, and you have an appropriate language for the dynamics of people offering transfers of resources and choices and obligations and so on. But ever since reading it I've always had a nagging suspicion that you should be able to somehow derive the dynamics of offering transfers of resources and choices and so on from substructural logics that we already understand well. Then the thinking real hard and writing down the right combinators might be reduced to 'that's just the way linear logic works'.

What I'm going to do below is *not that* — I'm not actually going to derive their combinators from anything, I don't know how to do that yet. I just took the idea of 'buyers and sellers of contracts' (which is to say, 'various parties making bets with one another') and tried to think about what shape it took in substructural logics, and I'm just going to disclose what seemed to work, and where that led me.

2 First Pass at Buyers and Sellers

2.1 Logical Preliminaries

I'm going to assume you already know well what inutitionistic (affine) linear logic is, and that I don't need to spell out the language — that's what I'm going to be working in, and I'll just sort of play fast and loose with what exactly is in the logic. I could have tried doing a classical logic, and in some ways that might have worked out more nicely, in that the price of two contracts, one that pays a dollar when proposition A is true, and another that pays a dollar when A is false, would have been neatly a dollar, by excluded middle. But I prefer to work in a constructive logic to see exactly "how much excluded middle" I'm using.

The first thing I'll need to definitely have in the language is a linear proposition that represents amounts money, and I want it to be arbitrarily divisible. So I'll say a hypothesis $\$(x)$ sitting in the context represents my posession of x

dollars, where x is an arbitrary nonnegative real number. Just assume (without further discussion of what mechanism enables it to be so) that things like

$$\$(x + y) \dashv\vdash \$(x) \otimes \$(y)$$

$$\$(0) \dashv\vdash 1$$

hold. Maybe there's an implicit context full of global rules like

$$!(\forall xy.\$(x + y) \multimap \$(x) \otimes \$(y))$$

or maybe $ isn't an atom but rather a defined proposition ginned up specifically to have those properties — or still something else, whatever. I don't care.

I'm going to treat the logic as affine (meaning you're allowed to throw resources away, i.e. you can prove $A \vdash 1$ but not necessarily $A \vdash A \otimes A$) so that *provable sequents* correspond to (maybe zero) *arbitrage opportunities*, or more blandly, to *transactions that every party consents to.* The interests of counterparties other than the 'me' that is the prover are represented by populating the context with various propositions, and my own interest in tolerating increases but not decreases in money comes exactly from affineness.

2.2 Definitions

If we're going to have bets — ahem, contracts — about 'real-world events', then we need a connection between (not-linear-resource, unrestricted, persistent) propositions and resources. I'm going to use the unrestricted implication $\Rightarrow$ for that. Make the abbreviation

$$[A] := A \Rightarrow \$1$$

To have a proposition $[A]$ in the context means I hold a contract that pays out a dollar whenever A is true. Just because the proposition A is unrestricted doesn't mean $A \Rightarrow \$1$ is; the contract is itself a linear resource. Its subject A can be proven true over and over, but you only get as many dollars for its truth as you have linear copies of the contract in your context..

When someone is willing to sell me a contract that pays out when A is true, that looks like

$$\mathsf{S}(A, p) := \$p \multimap [A]$$

When someone is willing to buy from me a contract that pays out when A is true, that looks like

$$\mathsf{B}(A, p) := [A] \multimap \$p$$

2.3 Consequences

Immediately we can prove some pretty obvious things like profiting off of market-making

$$\$p, \mathsf{S}(A, p), \mathsf{B}(A, p + q) \vdash \$(p + q)$$

An important thing the logic is keeping track of for us is that we must already have the $\$p$ in order to buy the contract from the seller, before we can resell it to the buyer.

Another simple provable thing is

$$\mathsf{S}(A,p), !(B \Rightarrow A) \vdash S(B,p)$$

If someone's willing to sell for $\$p$ a ticket that pays when A is true, and we know that B guarantees A being true, then it's rational for us to *act* as a seller of a B-contract at price $\$p$. Conceivably B might fail to hold while A is true for some other reason, in which case we would make a profit.

We can also prove some slightly less obvious things like

$$\mathsf{S}(A,p), \mathsf{S}(B,q) \vdash \mathsf{S}(A \vee B, p+q)$$

If we can buy an A-contract and a B-contract, then we can safely offer a contract (with a price that covers the cost of materials) that pays out when at least one of A and B is true — we will never have to pay out except under a circumstance in which at least one of the contracts we bought pays out.

An interesting special case of this example is when the propositions A and B are believed to exhaust the space, i.e. when $\vdash A \vee B$. In this case we would have

$$\mathsf{S}(A,p), \mathsf{S}(B,q) \vdash \mathsf{S}(\top, p+q)$$

For example, when A is something like 'team A wins the sportsball championship' and B is 'team B wins the sportsball championship' and no ties are allowed. If you looked at the betting markets and you saw the sum of the price of A and the price of B being less than \$1, you'd have an arbitrage opportunity, namely $S(\top, p+q)$, which is the market effectively selling you a 100% guaranteed \$1-generating ticket at the price of $\$(p+q) < \$1$.

2.4 Doubts about Time

At this point I notice there's an interesting and very basic problem-modelling question, that I'm going to state but defer answering conclusively for now. Assuming I 'know' $\vdash A \vee B$, and I use that fact in proving

$$A \vee B, \mathsf{S}(A,p), \mathsf{S}(B,q) \vdash \mathsf{S}(\top, p+q)$$

it feels like it might be cheating to do $\vee L$ on that assumption if which branch of the disjunction I wind up in affects my decision of what contracts to buy from sellers. What I really want a sequent proof to correspond to is a proof that I can make a transaction *right now* knowing only what I know; so maybe the $A \vee B$ needs to be guarded by some kind of temporal logic modality, that I know $A \vee B$ holds at some point in the future. And then perhaps the definition of $[A]$ would get modified as well, to be something like $(\Diamond A) \Rightarrow \$1$ or $(\bigcirc A) \Rightarrow \$1$ or something like that, depending on what modal logic you're in. Not sure yet. This feels related to how disjunctions are funny in type systems for distributed programming, where doing an elimination naively means that you might need to 'psychically' know the value of the disjunction at a remote host.

2.5 Buyer/Seller Duality

Normally with a prediction market (and with a classical-logic mindset) you'd expect some kind of duality between contracts on a proposition and its negation. Perhaps something like

$$\mathsf{B}(A, p) \dashv\vdash \mathsf{S}(\neg A, 1 - p)$$

for any price $p \in (0, 1)$. This being because a willingness to buy an A-pays-\$1 contract should be the same thing as being willing to buy a A, $\neg A$ bundle for \$1, and then selling off the $\neg A$ half for $\$(1-p)$. But this isn't provable directly with the definitions above. What *is* provable is rather interesting. Let's generalize S and B to

$$\mathbf{S}(A, X) := X \multimap [A]$$

$$\mathbf{B}(A, X) := [A] \multimap X$$

In this case we can get away with

$$(A \land B \Rightarrow 0), \mathbf{B}(A, X \multimap \$1) \vdash \mathbf{S}(B, X)$$

$$(A \lor B), \mathbf{S}(B, X) \vdash \mathbf{B}(A, X \multimap \$1)$$

Here's sketches of how the proofs go:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{X \multimap \$1, X \vdash \$1} \quad \dfrac{\dfrac{\dfrac{}{A, B \vdash A \land B} \quad \dfrac{}{0 \vdash \$1}}{A \land B \Rightarrow 0, A, B \vdash \$1}}{A \land B \Rightarrow 0, B \vdash A \Rightarrow \$1}}{A \land B \Rightarrow 0, (A \Rightarrow \$1) \multimap X \multimap \$1, X, B \vdash \$1}}{A \land B \Rightarrow 0, (A \Rightarrow \$1) \multimap X \multimap \$1 \vdash X \multimap (B \Rightarrow \$1)}}{A \land B \Rightarrow 0, \mathbf{B}(A, X \multimap \$1) \vdash \mathbf{S}(B, X)}$$

$$\dfrac{\dfrac{\dfrac{\dfrac{}{X \vdash X} \quad \dfrac{}{A \lor B, (B \Rightarrow \$1), (A \Rightarrow \$1) \vdash \$1}}{A \lor B, X \multimap (B \Rightarrow \$1), (A \Rightarrow \$1), X \vdash \$1}}{A \lor B, X \multimap (B \Rightarrow \$1) \vdash (A \Rightarrow \$1) \multimap X \multimap \$1}}{A \lor B, \mathbf{S}(B, X) \vdash \mathbf{B}(A, X \multimap \$1)}$$

So *if* A and B are contradictory and exhaustive, as A and $\neg A$ would be in classical logic, we have the equivalence of $\mathbf{S}(B, X)$ and $\mathbf{B}(A, X \multimap \$1)$. But $\$p \multimap \$1$ isn't quite the same thing as $\$(1 - p)$. What's going on there? The logic is again forcing us to keep track of the fact of how much initial capital we have to put up in order to exploit arbitrage opportunities.

The proposition $\$p \multimap \$1$ is an opportunity to profit by $\$(1 - p)$, but it's a weaker proposition than $\$(1 - p)$, precisely because you need to put up $\$p$ to get it going. We have

$$\$(1 - p) \vdash \$p \multimap \$1$$

but not the converse

$$\$p \multimap \$1 \nvdash \$(1-p)$$

A question I asked myself at this point was, to what extent can we get back some of the symmetry of the classical situation? Does there exist, perhaps, a proposition $Y(p)$ parameterized by a price p such that both

$$\$p \vdash Y(p)$$

$$Y(p) \dashv\vdash Y(1-p) \multimap \$1$$

It would be a weakening of the propositional meaning of $\$p$; something like 'you can make a profit of $\$p$ after perhaps some effort or initial capital', and it would be such a suitable weakening as to be a fixedpoint under the flip embodied by the second required axiom above.

Notably one attempt at a definition,

$$Y(p) := \$(1-p) \multimap \$1$$

does *not* work. We get

$$\$p \vdash \$(1-p) \multimap \$1$$

and

$$(\$p \multimap \$1) \multimap \$1 \vdash \$(1-p) \multimap \$1$$

is easy to prove, but

$$\$(1-p) \multimap \$1 \nvdash (\$p \multimap \$1) \multimap \$1$$

fails. After doing $\multimap R$, we can't proceed without any unconditional \$ hypotheses in the context.

I think it's possible to solve this puzzle, however, with a corecursively defined proposition:

$$Y(p) = \nu\alpha.\$p \oplus ((\$(1-p) \oplus (\alpha \multimap \$1)) \multimap \$1)$$

Here $\$p \multimap Y(p)$ is easy by left-injecting into the coproduct, and the other two proof obligations work by coinduction. First we prove a lemma:

$$\mathcal{D} = \dfrac{\dfrac{}{\$p, \$(1-p) \vdash \$1} \qquad \dfrac{\dfrac{\dfrac{}{\$p \vdash \$p}}{\$p \vdash \$p \oplus (Y(1-p) \multimap \$1)}}{\$p, (\$p \oplus (Y(1-p) \multimap \$1)) \multimap \$1 \vdash \$1}}{\$p, Y(1-p) \vdash \$1}\,\oplus L$$

and then we can show

$$
\dfrac{\dfrac{\dfrac{\mathcal{D}}{\$p, Y(1-p) \vdash \$1} \qquad \dfrac{\dfrac{\dfrac{[\text{coind. hyp.}]}{Y(p), Y(1-p) \vdash \$1}}{Y(1-p) \vdash \$(1-p) \oplus (Y(p) \multimap \$1)}\ \oplus R, \multimap R}{(\$(1-p) \oplus (Y(p) \multimap \$1)) \multimap \$1, Y(1-p) \vdash \$1}\ \multimap L}{Y(p), Y(1-p) \vdash \$1}\ \oplus L}{Y(p) \vdash Y(1-p) \multimap \$1}
$$

and in the opposite direction

$$
\dfrac{\dfrac{\dfrac{\dfrac{}{\$(1-p) \vdash Y(1-p)}}{Y(1-p) \multimap \$1, \$(1-p) \vdash \$1} \qquad \dfrac{\dfrac{[\text{coind. hyp.}]}{Y(1-p) \multimap \$1, \vdash Y(p)}}{Y(1-p) \multimap \$1, Y(p) \multimap \$1 \vdash \$1}}{Y(1-p) \multimap \$1, \$(1-p) \oplus (Y(p) \multimap \$1) \vdash \$1}}{\dfrac{Y(1-p) \multimap \$1 \vdash (\$(1-p) \oplus (Y(p) \multimap \$1)) \multimap \$1}{Y(1-p) \multimap \$1 \vdash Y(p)}}
$$

Having done all that, we can finally say that if $\vdash A \wedge B \Rightarrow 0$ and $\vdash A \vee B$, then

$$\mathbf{B}(A, Y(p)) \dashv\vdash \mathbf{S}(B, Y(1-p))$$

3 Definite Problems with Time

This section is a continuation of exploring the worry expressed in 2.4.

3.1 A Thing I Don't Want to Prove

Let's generalize the definitions so hard that they become just a trivially different (albeit suggestive) way of writing implications:

Say $[A]^p$ means $A \Rightarrow \$p$, pronounced "a ticket that pays out p when A is true".

Say $\mathbf{S}(Y, X) = X \multimap Y$, pronounced "there is a seller in the market willing to sell me contract Y at price X".

Say $\mathbf{B}(Y, X) = Y \multimap X$, pronounced "there is a seller in the market willing to by contract Y from me at price X".

In this case, there is a proposition I *can* prove, but shouldn't be able to according to my informal intended interpretation of the semantics. It is

$$\$p, \mathbf{B}([A]^q, \$m), \mathbf{B}([B]^t, \$q), \mathbf{S}([A \wedge B]^t, \$p) \vdash \$m$$

The proof is this:

$$\dfrac{\dfrac{}{\$m \vdash \$m} \qquad \dfrac{\dfrac{\dfrac{}{\$q \vdash \$q} \qquad \dfrac{\dfrac{\dfrac{}{\$p \vdash \$p} \qquad \dfrac{\dfrac{}{\$t \vdash \$t} \qquad \dfrac{}{A, B \vdash A \land B}}{A, B, [A \land B]^t \vdash \$t}}{A, B, \$p, \mathbf{S}([A \land B]^t, \$p) \vdash \$t}}{A, \$p, \mathbf{S}([A \land B]^t, \$p) \vdash [B]^t}}{A, \$p, \mathbf{B}([B]^t, \$q), \mathbf{S}([A \land B]^t, \$p) \vdash \$q}}{\$p, \mathbf{B}([B]^t, \$q), \mathbf{S}([A \land B]^t, \$p) \vdash [A]^q}}{\$p, \mathbf{B}([A]^q, \$m), \mathbf{B}([B]^t, \$q), \mathbf{S}([A \land B]^t, \$p) \vdash \$m}$$

The proof is a story of the transactions made: we sell a A-pays-q contract to the first buyer (call her Alice) in the context. This guarantees the revenue $\$m$ that the sequent says we expect to make at the end of the day. Now we need to show that we can definitely pay out the value of the contract we sold, under the assumption that A is true. Well, we can sell a B-pays-t ticket to Bob, at price q. That matches up with the $\$q$ we owe alice in case A happens. Now we must show that we can pay Bob t assuming B happens. Well by now we have both A and B in the context, so we can buy an $A \land B$-pays-t ticket from Carol with the initial capital $\$p$ that we assumed we had, and satisfy our obligation to Bob.

What's wrong with this story? It seems like a bad model to assume that a buyer at a cetain price that exists now will exist forever if *we* don't interact with them. So we actually want to imagine buying/selling all the contracts implicated in a single proof in the present, prior waiting for the truth of propositions of A and B to come to light.

So let's imagine that *right now* I sell the A-pays-q to Alice for $\$m$, I sell the B-pays-$\$t$ to Bob for $\$q$, and I buy a $A \land B$-pays-$\$t$ from Carol for $\$p$. I have $\$(m + q)$. If A comes true, I pay Alice $\$q$. If B then comes true, Carol pays me $\$t$ and I pass that along to Bob, and I'm left with at least $\$m$ in any case. So everything's fine, yeah?

Not so fast! What if B is known to be true *before* A? (or indeed B may be true and A false; but I'm being a good intuitionist and simply picturing the scenario where B is known to be true, and A isn't *known to be true yet*) In this case I have $\$(m+q)$ and I owe Bob $\$t$, and Carol doesn't owe me anything, and maybe $\$t$ is bigger than $\$(m + q)$ and I'm out of luck.

What went wrong here is I'm not being careful enough about the temporal character of the propositions becoming true and creating payment obligations. If $[A]^t$ was defined as somthing like $\Diamond A \Rightarrow \$t$, then instead of needing to prove $A, B \vdash A \land B$, we would have to prove $\Diamond A, \Diamond B \vdash \Diamond(A \land B)$, which is indeed *unprovable*, as desired! At least for most sensible interpretations of $\Diamond$ I have in mind. If A is true in some future world, and B is true in some future world, then you don't know it's the same world they're true in.

3.2 A Thing I Want to Prove

Do we then want a ‘plain old ordinary $\Diamond$’, like in [PD01], which gives you nothing particularly useful from $\Diamond A \wedge \Diamond B$? (At least nothing you couldn't get from $\Diamond A$ or $\Diamond B$ alone) Here's an argument why not.

Note that the bad case above is bad because $\$t$ *might be* bigger than $\$(m+q)$. If I choose all the numbers right so that they work for *either* ordering of the component events, then things should be ok.

I *do* think I should be able to prove, for example,

$$\mathbf{B}([A], \$(1-a)), \mathbf{B}([B], \$(1-b)), \mathbf{S}([A \wedge B], \$(1-a-b)) \vdash 1$$

for suitable definitions of all the propositions involved. Here I sell an A-pays-\$1 ticket to Alice, a B-pays-\$1 ticket to Bob, and buy an $A \wedge B$-pays-\$1 from Carol, leaving me with exactly $\$(1-a) + \$(1-b) - \$(1-a-b) = \$1$. If A comes true, I can pay that \$1 to Alice. If B comes true, then to Bob I forward Carol's payment of \$1. The same story plays out symmetrically if B is true first, and then A. If neither A or B are ever true, then I've made a profit.

So it seems like I'm going to want a modality that actually enables me as a prover to do case analysis over these orderings. Although — even if I have such a thing, I suspect $[A]^t = \Diamond A \Rightarrow \$t$ may still be too simplistic a definition of an A-pays-t contract. Even though the above scheme is self-funding, requiring no initial capital, the proof theory wouldn't let me apply the $\multimap L$ rule to $\mathbf{B}([A], \$(1-a))$ without throwing away $\$(1-a)$ in the $\$(1-a) \vdash 1$ branch.

3.3 Linear-time Possibility

But I suspect I do in any case want some kind of diamond that lets me prove

$$\Diamond A, \Diamond B \vdash \Diamond((\Diamond A \wedge B) \vee (A \wedge \Diamond B))$$

which is a standard thing to expect of $\Diamond$ when the Kripke relation is known to be a linear order. This behavior, in specifically the context of a constructive logic, i.e. a type theory, has a rather arrestingly pleasant application to FRP; Paykin and collaborators [PKZ15] add a variant of the above to their type theory as an axiom, where it is effectively the type of the unix system call `select`.

Instead of just throwing it in as an axiom, I have a burning desire to understand what kind of judgmental machinery could give rise to such a beast — I'm going to try to code it up out of ‘simpler’ parts, so that, e.g. the fact of its cut elimination falls out of the cut-well-behavedness of the parts. I say ‘simpler’ because the parts I'm going to build it out of are themselves slightly exotic as logical connectives go, and which I will treat very handwavily — but I think the decomposition is still perspicuous and interesting.

My proposal is the following:

$$\Diamond A = \exists x.[x] \otimes \Box(\forall y \geq x.[y] \multimap (A \vee [\bigcirc y]))$$

The new syntax I need to explain consists of $[\text{—}]$ and $\geq$ and $\bigcirc$.[1] The expression $[\text{—}]$ is an atomic proposition of the sort you'd expect to find in any discussion of first-order logic, whose argument is a term expression. Terms are built from term variables, the binary function symbol $*$, (which hasn't appeared yet, but will be required imminently) and the unary function symbol $\bigcirc$. There is a binary relation $\geq$ on terms.

I'll demand as axioms that $\leq$ is reflexive and transitive, and that

$$[x * y] \dashv\vdash [x] \otimes [y]$$

$$x \leq \bigcirc x \qquad \bigcirc(x * y) = \bigcirc x * \bigcirc y$$

$$\frac{x_1 \leq y_1 \qquad x_2 \leq y_2}{x_1 * x_2 \leq y_1 * y_2}$$

and furthermore whenever $y \geq x_1 * x_2$ then there exist $y_1 \geq x_1$ and $y_2 \geq x_2$ such that we can decompose $y = y_1 * y_2$.

To see how this definition works, let's go through the proof of

$$\Diamond A, \Diamond B \vdash \Diamond((\Diamond A \wedge B) \vee (A \wedge \Diamond B))$$

The first thing we do is unpack the existentials and tensors on the left. From there we have to prove:

$$\forall y \geq x_1.[y] \multimap (A \vee [\bigcirc y]) \ \mathsf{valid}, \forall y \geq x_2.[y] \multimap (B \vee [\bigcirc y]) \ \mathsf{valid}$$

$$[x_1], [x_2] \vdash \Diamond((\Diamond A \wedge B) \vee (A \wedge \Diamond B))$$

The $[x_1]$ and $[x_2]$ (together with those big valid assumptions that give them content) are kind of like A poss and B poss hanging out on the left of the turnstile. Now we work on the $\Diamond$ on the right: for the $\Box$ (which also has a ! baked in, remember?) to succeed, we have to choose something for x in the existential to gobble up the whole linear context. But choosing $x = x_1 * x_2$ does exactly this, since then $[x] = [x_1 * x_2] = [x_1] \otimes [x_2]$. So now our proof obligation is: (refraining from writing down again the valid assumptions, which are still there)

$$\vdash \forall y \geq (x_1 * x_2).[y] \multimap ((\Diamond A \wedge B) \vee (A \wedge \Diamond B)) \vee [\bigcirc y]$$

Use the axiom about decomposing $y \geq x_1 * x_2$ to see this is the same as

$$\vdash \forall y_1 \geq x_1.\forall y_2 \geq x_2.[y_1 * y_2] \multimap ((\Diamond A \wedge B) \vee (A \wedge \Diamond B)) \vee [\bigcirc(y_1 * y_2)]$$

and then let all the asynchronous stuff fire, and turn the crank on the axioms some more, and we get

$$[y_1], [y_2] \vdash ((\Diamond A \wedge B) \vee (A \wedge \Diamond B)) \vee ([\bigcirc y_1] \otimes [\bigcirc y_2])$$

[1] The $\Box$ is pretty much what you'd expect of $\Box$ in intuitionistic logic, (see [PD01] for what *I* expect, anyhow) except take my word for it that it also has ! baked into it; it assumes (on the left, or requires on the right) that everything is valid not merely true, and valid things are definitely persistent not linear.

Where can we go from here? Focusing on the right probably won't work. The left branch of the $\vee$ is probably a non-starter since any As and Bs in our hypotheses are still locked up in the valid context. We can't prove $[\bigcirc y_1] \otimes [\bigcirc y_2]$ either, since all we have is $[\bigcirc y_1] \otimes [\bigcirc y_2]$. So let's try actually pushing the button on one of the valid assumptions. Use $[y_1]$ (since we know $y_1 \geq x_1$) together with $\forall y \geq x_1.[y] \multimap (A \vee [\bigcirc y])\ \mathsf{valid}$ to work up to

$$\dfrac{\dfrac{A, [y_2] \vdash C \qquad [\bigcirc y_1], [y_2] \vdash C}{A \vee [\bigcirc y_1], [y_2] \vdash C}}{[y_1], [y_2] \vdash C}$$

(where $C = ((\Diamond A \wedge B) \vee (A \wedge \Diamond B)) \vee ([\bigcirc y_1] \otimes [\bigcirc y_2])$) This is starting to look good. In the left branch, we might have a hope of proving $A \wedge \Diamond B$, and in the right branch, we're not far from proving $[\bigcirc y_1] \otimes [\bigcirc y_2]$. We do it by sending y_2 into $\forall y \geq x_2.[y] \multimap (B \vee [\bigcirc y])\ \mathsf{valid}$:

$$\dfrac{\dfrac{A, [y_2] \vdash C \qquad \dfrac{\dfrac{[\bigcirc y_1], B \vdash C \qquad \dfrac{\dfrac{}{[\bigcirc y_1], [\bigcirc y_2] \vdash [\bigcirc y_1] \otimes [\bigcirc y_2]}}{[\bigcirc y_1], [\bigcirc y_2] \vdash C}}{[\bigcirc y_1], B \vee [\bigcirc y_2] \vdash C}}{[\bigcirc y_1], [y_2] \vdash C}}{A \vee [\bigcirc y_1], [y_2] \vdash C}}{[y_1], [y_2] \vdash C}$$

The informal story going on here is that we have two events, A and B; we waiting for A and maybe it happened or not. If A didn't happen yet, we waited for B and it happened or not. If B didn't happen either, then we learn that *none of the events we were waiting for happened*, and this is a sufficient observation to terminate the proof.

In the remaining branches, we have some event that definitely did happen, and all of the remaining events are still waiting to happen. If we can show $A, [y_2] \vdash A \wedge \Diamond B$ and $B, [\bigcirc y_1] \vdash \Diamond A \wedge B$ and we're done. The apparent asymmetry between $[y_2]$ and $[\bigcirc y_1]$ is smoothed over by the fact that the $\forall$ in the definition of $\Diamond$ is over any $y \geq x$; it doesn't care if there's a few $\bigcirc$s tacked on: we assumed $x \leq \bigcirc x$ for any x, and that $\geq$ is transitive. So the proof that $[\bigcirc y_1] \vdash \Diamond A$ looks like

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{A \vee [\bigcirc y] \vdash A \vee [\bigcirc y]}\,{*}}{y \geq \bigcirc y_1, [y] \vdash A \vee [\bigcirc y]}}{\vdash \Box(\forall y \geq \bigcirc y_1.[y] \multimap (A \vee [\bigcirc y]))}}{[\bigcirc y_1] \vdash [\bigcirc y_1] \otimes \Box(\forall y \geq \bigcirc y_1.[y] \multimap (A \vee [\bigcirc y]))}}{[\bigcirc y_1] \vdash \exists x.[x] \otimes \Box(\forall y \geq x.[y] \multimap (A \vee [\bigcirc y]))}$$

where $*$ is the moment where we apply $\forall y \geq x_1.[y] \multimap (A \vee [\bigcirc y])$ valid. The proof of $[y_2] \vdash \Diamond B$ is virtually identical, except with one fewer $\bigcirc$s in play.

Some parting thoughts on this proposal, as informal as it is:

1. From the point of view of focusing/polarization, it's a huge beast! Three phases! Positive, negative, positive. I don't think I've seen that before in any sensible connective. Is there a way of simplifying it to two somehow?

2. It seems rather more $\Box$-ish than $\Diamond$-ish. There's a $\Box$ right in the definition for heaven's sake, and it's not dualized or anything. What's up with that? Is there a dual version of it that means 'for all future times on a linear timeline' that seems equally $\Diamond$-ish, perhaps?

3. It doesn't satisfy $A \vdash \Diamond A$, but you could probably get around that by just considering $A \vee \Diamond A$.

4. It doesn't satisfy $\Diamond\Diamond A \vdash \Diamond A$, and I don't know how to easily fix that.

5. It doesn't satisfy $\vdash \Diamond 0$, which is good, and that's exactly what the $\bigcirc$ is there for. Take it out and you do get $\vdash \Diamond 0$.

6. It doesn't satisfy $\Diamond 0 \vdash 0$, which I tend to think is proper. But it seems kind of fragile, like it's dangerously *almost* provable. You can still turn the crank on the $\Diamond 0$ assumption, but all you do is run up one branch of the proof tree getting $[\bigcirc y]$ then $[\bigcirc\bigcirc y]$ then $[\bigcirc\bigcirc\bigcirc y]$ etc. and you never have a $\Diamond$ on the right to connect it up with.

References

[PD01] Frank Pfenning and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical. Structures in Comp. Sci.*, 11(4):511–540, August 2001.

[PJES00] Simon Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: An adventure in financial engineering (functional pearl). *SIGPLAN Not.*, 35(9):280–292, September 2000.

[PKZ15] Jennifer Paykin, Neelakantan R Krishnaswami, and Steve Zdancewic. Linear temporal type theory for event-based reactive programming, 2015.